

A large deviation principle for join the shortest queue

Anatolii A. Puhalskii

University of Colorado at Denver, Denver, U.S.A. and
Institute of Information Transmission Problems, Moscow, Russia

Alexander A. Vladimirov

Institute of Information Transmission Problems, Moscow, Russia

February 2, 2008

Abstract

We consider a join-the-shortest-queue model which is as follows. There are K single FIFO servers and M arrival processes. The customers from a given arrival process can be served only by servers from a certain subset of all servers. The actual destination is the server with the smallest weighted queue length. The arrival processes are assumed to obey a large deviation principle while the service is exponential. A large deviation principle is established for the queue-length process. The action functional is expressed in terms of solutions to mathematical programming problems. The large deviation limit point is identified as a weak solution to a system of idempotent equations. Uniqueness of the weak solution is proved by establishing trajectorial uniqueness.

1 Introduction and summary

Motivation Queueing models with the join-the-shortest-queue (abbreviated further as JSQ) mechanism for routing arriving customers are of interest in various application areas, see, e.g., Fleming and Simon [14], Turner [31]. In a generic set-up, one considers a queueing system consisting of service stations arranged in parallel. The customers arrive exogenously and join the station that has the least number of customers to serve. If there are several stations that have fewest customers, then some rule is applied to direct an arriving customer to one of these stations. The model admits many versions: along with the stream of the “discretionary” customers that choose the station with the least number of customers there may be “dedicated” customers who can only be served at specific stations, there may be several streams of discretionary customers each of which can join only the stations from a certain subset of the set of all stations, the number of servers at the service stations can vary from one to infinity, the stations may be assigned weights so that the discretionary customers join the station with the least weighted number of customers, upon service completion the customers may either leave the system or be routed back, and so on. However, all these versions share the common feature that the arrival rates at the stations depend discontinuously on the numbers of customers present. More precisely, the rate at which a given station receives customers experiences a jump when the set of the stations with fewest customers changes.

Keywords and phrases: join the shortest queue, large deviations, idempotent probability, discontinuous dynamics
MSC 2000 subject classifications: primary 60F10, secondary 60K25

Thus, JSQ models fall in the category of stochastic systems with discontinuous dynamics (or discontinuous stochastic dynamical systems). The dynamical properties of such a system change abruptly when it enters a certain domain of the state space. In this paper we concern ourselves with a study of the large deviation principle (LDP) for the queue lengths in a JSQ model. Large deviations of discontinuous stochastic dynamical systems have received considerable attention in the literature, see Alanyali and Hajek [2], Atar and Dupuis [3], Blinovski and Dobrushin [4], Borovkov and Mogulskii [5], Boué, Dupuis and Ellis [6], Dupuis and Ellis [9, 10, 11], Dupuis, Ellis and Weiss [12], Ignatiouk [17, 18], Korostelev and Leonov [20, 21], Majewski [23], Ramanan and Dupuis [28], Shwartz and Weiss [30]. However, available results on the LDP for JSQ are confined to the setting of two stations, see Ridder and Shwartz [29] and references therein. Even for that case we have not been able to find in the literature an explicit formulation of an LDP for the queue length process with an indication of the associated action functional. Paths to overflow have been studied in Foley and McDonald [15], Ridder and Shwartz [29], and Turner [31]. A version of the model where the stations are represented by infinite servers has been analysed by Alanyali and Hajek [1] who for the case of two stations establish an LDP for the queue length trajectories and study overflow paths, and Turner [31] who contrasts the results for the single-server and infinite-server cases.

Model We consider a JSQ model with an arbitrary number of arrival processes and an arbitrary number of service stations which are represented by single servers with the FIFO service discipline. The customers from a given arrival process can be served only by the servers from a certain subset of all servers. The actual destination is the server with the smallest weighted queue length. The model thus incorporates both dedicated and discretionary flows of customers. The arrival processes are general and are only assumed to obey an LDP. The service times are associated with the servers and are exponentially distributed. (In fact, our main result concerns a more general setting of autonomous service.) Upon service completion, the customers depart from the system.

Methods We adopt the strategy that has proved to be useful for establishing LDPs for continuous dynamical systems. It is based on the characterisation of large deviation relatively compact sequences as exponentially tight ones, see, e.g., Puhalskii [24], [26]. The large deviation limit point is identified as a weak solution to idempotent equations. The latter equations are obtained as large deviation limits of the stochastic equations governing the original dynamical system.

However, this general method has to be modified for stochastic systems with discontinuous dynamics, in particular, for the JSQ model we are concerned with here. Building on the approach of Puhalskii [27], we first replace the original discontinuous equations with certain continuous ones, to which the limit procedure is applied. We prove trajectory uniqueness for the resulting system of idempotent equations. This implies uniqueness of the large deviation limit point, hence, an LDP. The trajectory uniqueness is proved by showing the existence of a Lyapunov function.

Contribution We establish an LDP for the queue length process considered as a random element of the associated Skorohod space. The action functional is of integral form and is expressed in terms of solutions to mathematical programming problems. We also provide some insight into the relation between weak and trajectory uniqueness for idempotent equations which is instrumental in the proof. Besides, uniqueness of solutions for a fluid version of the JSQ model is established.

In broader perspective, the approaches of this paper and those of Puhalskii [27] provide new tools for the study of large deviations of stochastic systems with discontinuous dynamics. The system of idempotent equations that needs to have a weakly unique solution in order for an LDP to hold has the form of a system of differential equations. The latter system can be viewed as

describing a fluid version of the original queueing model where the arrival, service and queue length processes are represented by absolutely continuous functions. One of the implications of the results of this paper is that if, given trajectories of exogenous arrival and service processes (“the inputs”), the system of differential equations has a unique solution for the queue length trajectory (“the output”) then weak uniqueness for the associated idempotent equations holds, so an LDP holds. In short, trajectory uniqueness for the fluid model implies an LDP. In fact, one can allow the fluid model to have a unique solution for a certain subset of the set of all “inputs”, cf., Puhalskii [27]. As the application in this paper shows, the method does not restrict the number of domains of constant dynamics sharing a common boundary, while such constraints are inherent in the existing techniques.

Organisation The paper is organised as follows. The main result is stated in Section 2. Section 3 is concerned with its proof. The appendix reviews basics of idempotent probability and large deviation convergence.

Notation, terminology, conventions We will say that a function \mathbf{I} from a metric space \mathbb{E} to $[0, \infty]$ is an action functional if it is lower compact that is the sets $\{z \in \mathbb{E} : \mathbf{I}(z) \leq a\}$ are compact for $a \in \mathbb{R}_+$ and $\inf_{z \in \mathbb{E}} \mathbf{I}(z) = 0$. A sequence $\{\mathbf{P}_n, n \in \mathbb{N}\}$ of probability measures on the Borel σ -algebra of \mathbb{E} (or a sequence of random elements $\{X_n, n \in \mathbb{N}\}$ with values in \mathbb{E} and distributions \mathbf{P}_n) is said to obey the LDP for scale n with the action functional \mathbf{I} if $\limsup_{n \rightarrow \infty} n^{-1} \log \mathbf{P}_n(F) \leq -\inf_{z \in F} \mathbf{I}(z)$ for each closed subset F of \mathbb{E} and $\liminf_{n \rightarrow \infty} n^{-1} \log \mathbf{P}_n(G) \geq -\inf_{z \in G} \mathbf{I}(z)$ for each open subset G of \mathbb{E} .

The set of natural numbers is denoted as \mathbb{N} , the set of real numbers is denoted as \mathbb{R} and the non-negative halfline is denoted as \mathbb{R}_+ ; for $\ell \in \mathbb{N}$, \mathbb{R}^ℓ and \mathbb{R}_+^ℓ denote the cartesian products of ℓ copies of \mathbb{R} and \mathbb{R}_+ , respectively, with product topology. Elements of \mathbb{R}^ℓ are considered as column-vectors, 1_ℓ denotes the element of \mathbb{R}^ℓ with all the entries equal to one, superscript T is used to denote the transpose of a matrix. Inequalities involving vectors are understood componentwise. We denote as $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^\ell)$ the Skorohod space of \mathbb{R}^ℓ -valued right-continuous with left-hand limits functions on \mathbb{R}_+ . It is assumed to be endowed with the Skorohod J_1 -topology and metrised by a complete separable metric, see Ethier and Kurtz [13], Liptser and Shiryaev [22], Jacod and Shiryaev [19] for the definition and properties. The elements of $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^\ell)$ are denoted with lower-case bold-face Roman characters, e.g., $\mathbf{x} = (\mathbf{x}(t), t \in \mathbb{R}_+)$.

The following conventions are assumed: sums and infima over the empty set are equal to 0 and ∞ , respectively, $0/0 = 0$, and $0 \cdot (-\infty) = 0 \cdot \infty = \infty \cdot 0 = 0$. For $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}$, we denote as $\lfloor \alpha \rfloor$ the integer part, $\alpha \wedge \beta = \min(\alpha, \beta)$, $\alpha \vee \beta = \max(\alpha, \beta)$, and $\alpha^+ = \alpha \vee 0$; $\mathbf{1}(\Gamma)$ denotes the indicator function of a set Γ that is equal to 1 on Γ and is equal to 0 outside of Γ . We use dot to denote differentiation with respect to the time variable, so that $\dot{\mathbf{x}}(t)$ denotes the time-derivative of an absolutely continuous function $(\mathbf{x}(t), t \in \mathbb{R}_+)$; if $\mathbf{x}(t)$ is a vector, then the notation signifies that each component is differentiated; $\mathbf{x}(t-)$ denotes the left-hand limit at t ; “a.e.” refers to Lebesgue measure unless specified otherwise. All the relations involving derivatives are understood to hold a.e.

2 The LDP for join the shortest queue

Dynamics We provide a more specific model description and turn it into equations relating stochastic processes. There are K queues, indexed 1 through K , each with a single server. The servers are fed by M arrival processes indexed 1 through M . Customers from the m -th arrival

process can be served only by the servers from a nonempty subset S_m of all servers. More specifically, an arriving customer is routed to the queue with the smallest weighted queue length: queue $k \in S_m$ is assigned weight w_{km} , where $w_{km} > 0$, and the customer in question joins the queue with the least value of x_k/w_{km} , $k = 1, 2, \dots, K$, where x_k stands for the number of customers in queue k at the time when the arrival occurs. In case there are two or more queues with the least value of x_k/w_{km} , the arrival is routed to one of the competing queues. The actual rule used for breaking ties is of no consequence for the results below. Note that in this set-up dedicated arrivals are those with one-element sets S_m .

We assume that the servers perform service autonomously in the following sense. Each server is assigned a point process (i.e., a piecewise constant nondecreasing process starting at zero with unit jumps) which specifies service completion times: if at the moment preceding a jump time there is a customer present in the queue, then this customer leaves the queue at the time of the jump. Due to the memoryless property of the exponential distribution, servers with exponential service times can be considered as servers with autonomous service.

Our set-up concerns, in fact, a sequence of queueing models as described above indexed by $n \in \mathbb{N}$. To distinguish between the models, superscript n will be used in the notation for the entities associated with the n -th model. Let $A_m^n = (A_m^n(t), t \in \mathbb{R}_+)$ and $B_k^n = (B_k^n(t), t \in \mathbb{R}_+)$, where $m = 1, 2, \dots, M$ and $k = 1, \dots, K$, be one-dimensional point processes. It is assumed that $A_m^n(0) = B_k^n(0) = 0$. We let $A_m^n(t)$ model the cumulative number of exogenous arrivals in the m -th arrival process by time t and let $B_m^n(t)$ model the number of service completions by server m during time t of uninterrupted work of the server. All the processes are defined on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ and have trajectories in the associated Skorohod spaces. Let C_k denote the set of $m \in \{1, 2, \dots, M\}$ such that $k \in S_m$ and $Q_k^n(t)$ denote the number of customers in queue k at time t . The process $Q^n = (Q^n(t), t \in \mathbb{R}_+)$, where $Q^n(t) = (Q_k^n(t), k = 1, 2, \dots, K)$, satisfies the following balance equations, where $k = 1, 2, \dots, K$ and $t \in \mathbb{R}_+$,

$$\begin{aligned} Q_k^n(t) &= Q_k^n(0) + \sum_{m \in C_k} \left[\int_0^t \mathbf{1} \left(\frac{Q_k^n(s-)}{w_{km}} < \min_{\substack{l \in S_m: \\ l \neq k}} \frac{Q_l^n(s-)}{w_{lm}} \right) dA_m^n(s) \right. \\ &\quad \left. + \int_0^t \mathbf{1} \left(\frac{Q_k^n(s-)}{w_{km}} = \min_{\substack{l \in S_m: \\ l \neq k}} \frac{Q_l^n(s-)}{w_{lm}} \right) \alpha_{km}^n(s) dA_m^n(s) \right] - \int_0^t \mathbf{1}(Q_k^n(s-) > 0) dB_k^n(s). \end{aligned} \quad (2.1)$$

In this equation, the random variables $\alpha_{km}^n(s)$ account for the rule adopted for breaking ties between several queues of minimum weighted length, so one of the random variables $\alpha_{km}^n(s)$, $k = 1, 2, \dots, K$, equals 1 while the rest equal 0. The integrals on the right of (2.1) are well defined by being finite sums. Since the A_m^n and B_k^n have piecewise constant trajectories with finite numbers of jumps on bounded intervals, (2.1) admits a unique solution for $Q_k^n(t)$, see Chen and Mandelbaum [7] for more extensive results.

Statement of the LDP Let us be given a $[0, \infty]$ -valued Borel function $\psi(z)$, where $z \in \mathbb{R}_+^{M+K}$. Let function $\mathbf{I}^{A,B} : \mathbb{D}(\mathbb{R}_+, \mathbb{R}^{M+K}) \rightarrow \mathbb{R}_+$ be defined by the equality

$$\mathbf{I}^{A,B}(\mathbf{z}) = \int_0^\infty \psi(\dot{\mathbf{z}}(t)) dt, \quad (2.2)$$

if the function $\mathbf{z} = (\mathbf{z}(t), t \in \mathbb{R}_+)$ is absolutely continuous, componentwise nondecreasing, and $\mathbf{z}(0) = 0$, and $\mathbf{I}^{A,B}(\mathbf{z}) = \infty$ otherwise. It is assumed that $\mathbf{I}^{A,B}$ is an action functional on $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^{M+K})$. It follows that ψ is an action functional itself.

For $x = (x_1, \dots, x_K) \in \mathbb{R}_+^K$ and $y = (y_1, \dots, y_K) \in \mathbb{R}^K$, let $N(x, y)$ denote the set of $a = (a_1, \dots, a_M) \in \mathbb{R}_+^M$ and $b = (b_1, \dots, b_K) \in \mathbb{R}_+^K$ for which there exist matrices $e = (e_{km}) \in \mathbb{R}_+^{K \times M}$ and vectors $d = (d_1, \dots, d_K) \in \mathbb{R}_+^K$ such that $y = e1_M - d$, $e^T 1_K \leq a$, $d \leq b$, $e_{km} = 0$ if either $k \notin S_m$ or $x_k/w_{km} > \min_{l \in S_m} x_l/w_{lm}$, and $d_k = a_k$ if $x_k > 0$. Let

$$L(x, y) = \inf_{(a, b) \in N(x, y)} \psi(a, b). \quad (2.3)$$

Since ψ is lower compact and $N(x, y)$ is closed, the infimum above is attained provided $N(x, y)$ is nonempty. Besides, an easy compactness argument shows that the function L is lower semicontinuous. In particular, it is Borel measurable.

For $q_0 \in \mathbb{R}^K$ and $\mathbf{q} \in \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$, we define

$$\mathbf{I}_{q_0}^Q(\mathbf{q}) = \int_0^\infty L(\mathbf{q}(t), \dot{\mathbf{q}}(t)) dt$$

if $q_0 \in \mathbb{R}_+^K$, the function \mathbf{q} is componentwise nonnegative, absolutely continuous and $\mathbf{q}(0) = q_0$, and $\mathbf{I}_{q_0}^Q(\mathbf{q}) = \infty$ otherwise.

Let $\overline{Q}_k^n(t) = Q_k^n(nt)/n$, $\overline{A}_m^n(t) = A_m^n(nt)/n$, and $\overline{B}_k^n(t) = B_k^n(nt)/n$, where $m = 1, 2, \dots, M$ and $k = 1, 2, \dots, K$. We introduce the processes $\overline{Q}_k^n = (\overline{Q}_k^n(t), t \in \mathbb{R}_+)$, $\overline{A}_m^n = (\overline{A}_m^n(t), t \in \mathbb{R}_+)$, $\overline{B}_m^n = (\overline{B}_m^n(t), t \in \mathbb{R}_+)$, $\overline{Q}^n = (\overline{Q}_k^n, k = 1, 2, \dots, K)$, $\overline{A}^n = (\overline{A}_m^n, m = 1, 2, \dots, M)$, and $\overline{B}^n = (\overline{B}_k^n, k = 1, 2, \dots, K)$.

Theorem 2.1. *Let, as $n \rightarrow \infty$, the sequence $\{(\overline{Q}^n(0), \overline{A}^n, \overline{B}^n), n \in \mathbb{N}\}$ obey the LDP in $\mathbb{R}^K \times \mathbb{D}(\mathbb{R}_+, \mathbb{R}^M) \times \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$ for scale n with the action functional $\mathbf{I}^{Q_0, A, B}(q_0, \mathbf{a}, \mathbf{b}) = \mathbf{I}^{Q_0}(q_0) + \mathbf{I}^{A, B}(\mathbf{a}, \mathbf{b})$, where \mathbf{I}^{Q_0} is an action functional on \mathbb{R}^K . Then the sequence $\{\overline{Q}^n, n \in \mathbb{N}\}$ obeys the LDP for scale n in $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$ with the action functional \mathbf{I}^Q defined by the equality $\mathbf{I}^Q(\mathbf{q}) = \mathbf{I}^{Q_0}(\mathbf{q}(0)) + \mathbf{I}_{\mathbf{q}(0)}^Q(\mathbf{q})$.*

Comments and corollaries The fact that \mathbf{I}^Q is an action functional under the hypotheses is a part of the assertion of the theorem. One can establish this property without relying on the proof of Theorem 2.1 if in addition to being an action functional the function ψ is assumed to be convex and of superlinear growth at infinity. The latter conditions also ensure that $\mathbf{I}^{A, B}$ is an action functional. The argument is similar to the one in Puhalskii [27].

Expression (2.3) for the local action functional is intuitive. If we interpret a as the vector of instantaneous exogenous arrival rates, b , as the vector of service rates, d , as the vector of customer departure rates, and e , as the matrix of rates at which customers from different exogenous arrival processes arrive at the servers, then (2.3) tells us that “the cost” for the queue-length vector to change at rate y given it equals x is obtained by minimising the local action functional for the exogenous arrivals and service subject to certain “conservation laws”: the rate of the queue length change must equal the difference between the arrival and departure rates at the servers, the departure rates cannot exceed the service rates, the sum over all the servers of the arrival rates due to the customers originating from a specific exogenous arrival process must equal the arrival rate of this process. Besides, the ratio d_k/b_k , provided $b_k > 0$, can be interpreted as the fraction of time that server k is busy serving customers, while e_{km}/a_m can be interpreted as the fraction of stream m customers directed to server k .

The definition of the function $L(x, y)$ implies that it is piecewise constant in x in that

$$L(x, y) = \sum_{I, J} \mathbf{1}(x \in F_{IJ}) \Psi_{IJ}(y). \quad (2.4)$$

The summation is over all subsets I of $\{1, 2, \dots, K\}$ (including the empty set) and all sets J of the form $J = J_1 \times \dots \times J_M$, where the J_m , for $m = 1, 2, \dots, M$, are nonempty subsets of the S_m such that either $J_m \subset I$ or $J_m \cap I = \emptyset$, F_{IJ} denotes the subset of \mathbb{R}_+^K of elements $x = (x_1, \dots, x_K)$ such that $x_k = 0$ if $k \in I$, $x_k > 0$ if $k \notin I$, $x_k/w_{km} = \min_{l \in S_m} x_l/w_{lm}$ for $m = 1, 2, \dots, M$ if $k \in J_m$, and $x_k/w_{km} > \min_{l \in S_m} x_l/w_{lm}$ for $m = 1, 2, \dots, M$ if $k \notin J_m$. In words, if $x \in F_{IJ}$, then set I indexes the entries of x that are equal to zero and J_m indexes the entries that are the smallest weighted fluid queue lengths for the m -th fluid arrival process. The functions Ψ_{IJ} are uniquely specified by (2.4). It is easy to see that these functions are action functionals on \mathbb{R}^K . If, moreover, the function ψ is convex, then the functions Ψ_{IJ} are also convex. Representation (2.4) shows that the F_{IJ} are the domains of constant queue length dynamics.

The functions Ψ_{IJ} can be written down more explicitly if more structure is imposed on the function ψ . If we assume that $\psi(a, b) = \psi^A(a) + \sum_{k=1}^K \psi_k^B(b_k)$, where the functions $\psi_k^B : \mathbb{R}_+ \rightarrow [0, \infty]$ are lower semicontinuous, convex, attain zero, and are not equal to zero identically, then minimisation with respect to b in the definition of Ψ_{IJ} yields the representation

$$\Psi_{IJ}(y) = \inf_{(a, d) \in H_J^{-1}(y)} \left(\psi^A(a) + \sum_{k \notin I} \psi_k^B(d_k) + \sum_{k \in I} \psi_k^B(d_k) \mathbf{1}(d_k > \mu_k) \right),$$

where $\mu_k = \sup\{b_k \in \mathbb{R}_+ : \psi_k^B(b_k) = 0\}$ and $H_J(a, d)$, for $a \in \mathbb{R}_+^M$ and $d \in \mathbb{R}_+^K$, is the set of $\tilde{y} \in \mathbb{R}^K$ for which there exist matrices $e = (e_{km}) \in \mathbb{R}_+^{K \times M}$ such that $\tilde{y} = e1_M - d$, $e^T 1_K \leq a$, and $e_{km} = 0$ if $k \notin J_m$.

As a consequence of this observation and Theorem 2.1, we obtain the following result for the Markovian setting. For $\alpha \in \mathbb{R}_+$, we denote $\pi(\alpha) = \alpha \log \alpha - \alpha + 1$.

Theorem 2.2. *Let the A_m^n and B_k^n be independent Poisson processes with respective rates λ_m^n and μ_k^n , which are also independent of $Q^n(0)$. Let, as $n \rightarrow \infty$, $\lambda_m^n \rightarrow \lambda_m$, $\mu_k^n \rightarrow \mu_k$, and the sequence $\{\overline{Q}^n(0), n \in \mathbb{N}\}$ obey the LDP for scale n in \mathbb{R}^K with an action functional \mathbf{I}^{Q_0} . Then the sequence $\{\overline{Q}^n, n \in \mathbb{N}\}$ obeys the LDP for scale n in $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$ with the action functional $\mathbf{I}^Q(\mathbf{q}) = \mathbf{I}^{Q_0}(\mathbf{q}(0)) + \mathbf{I}_{\mathbf{q}(0)}^Q(\mathbf{q})$. The function $L(x, y)$ is of the form (2.4), where*

$$\Psi_{IJ}(y) = \inf_{(a, d) \in H_J^{-1}(y)} \left(\sum_{m=1}^M \pi\left(\frac{a_m}{\lambda_m}\right) \lambda_m + \sum_{k \notin I} \pi\left(\frac{d_k}{\mu_k}\right) \mu_k + \sum_{k \in I} \pi\left(\frac{d_k}{\mu_k}\right) \mu_k \mathbf{1}\left(\frac{d_k}{\mu_k} > 1\right) \right).$$

For the proof of Theorem 2.2, note that by Theorem 2.3 in Puhalskii [25] the LDP for $\{(\overline{Q}^n(0), \overline{A}^n, \overline{B}^n), n \in \mathbb{N}\}$ in the hypotheses of Theorem 2.1 holds with $\psi(a, b) = \sum_{m=1}^M \pi(a_m/\lambda_m) \lambda_m + \sum_{k=1}^K \pi(b_k/\mu_k) \mu_k$.

3 Proof of Theorem 2.1

The proof uses the terminology of large deviation convergence (abbreviated as LD convergence) and idempotent probability, which is recapitulated in the appendix, for more detail see Puhalskii [26]. We start with a proof outline, where we also define some concepts which illuminate the connection of LD convergence and weak convergence.

Weak uniqueness and trajectorial uniqueness Let $X^n = (\overline{Q}^n(0), \overline{A}^n, \overline{B}^n)$ and $Y^n = \overline{Q}^n$. These random variables assume values in the respective metric spaces $\mathbb{E}^X = \mathbb{R}^K \times \mathbb{D}(\mathbb{R}_+, \mathbb{R}^M) \times \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$ and $\mathbb{E}^Y = \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$. The sequence X^n LD converges in distribution at rate n by hypotheses. We seek to prove that the Y^n LD converge in distribution and find the limit.

Equation (2.1) is transformed into an equation $F_t(X^n, Y^n) = 0$, $t \in \mathbb{R}_+$, where the functions F_t are discontinuous. This equation enables us to establish that the sequence Y^n is exponentially tight. Thus, there exists a subsequence $(X^{n'}, Y^{n'})$ that LD converges in distribution to a certain idempotent process (X, Y) . By hypotheses, the idempotent process X has idempotent distribution $\mathbf{\Pi}^X(\mathbf{x}) = \exp(-\mathbf{I}^{Q_0, A, B}(\mathbf{x}))$, where $\mathbf{x} \in \mathbb{E}^X$. In order to identify the idempotent distribution of Y we would like to relate X and Y by equations obtained as certain LD limits. Since the original equation $F_t(X^n, Y^n) = 0$, $t \in \mathbb{R}_+$ involves discontinuities, as a preliminary step we replace it with an equation $\hat{F}_t(X^n, Y^n, Z^n) = 0$, $t \in \mathbb{R}_+$, where the \hat{F}_t are continuous functions and Z^n are additional random variables, which we call ‘‘latent’’. They assume values in a metric space \mathbb{E}^Z . The sequence (X^n, Y^n, Z^n) is still provably exponentially tight. Taking an LD limit along a subsequence, we have that if idempotent variables (X, Y, Z) defined on an idempotent probability space $(\Upsilon, \mathbf{\Pi})$ constitute an LD accumulation point of the (X^n, Y^n, Z^n) for LD convergence in distribution, then $\hat{F}_t(X, Y, Z) = 0$ $\mathbf{\Pi}$ -a.e. (Note that one can always take $\mathbb{E}^X \times \mathbb{E}^Y \times \mathbb{E}^Z$ as Υ .) The idempotent distribution of (X, Y, Z) is proved to be concentrated on a set $\mathbb{E}_0 \subset \mathbb{E}^X \times \mathbb{E}^Y \times \mathbb{E}^Z$ in the sense that $\mathbf{\Pi}((X, Y, Z) \in \mathbb{E}^X \times \mathbb{E}^Y \times \mathbb{E}^Z \setminus \mathbb{E}_0) = 0$. We refer to the idempotent distribution of Y , which is defined by $\mathbf{\Pi}^Y(\mathbf{y}) = \mathbf{\Pi}(Y = \mathbf{y})$, as a weak solution of the equation $\hat{F}_t(X, Y, Z) = 0$. If this weak solution is unique, then the idempotent law of Y is specified uniquely, so the Y^n LD converge in distribution to Y . As in the theory of stochastic differential equations, weak uniqueness follows from trajectorial uniqueness, which is defined as follows. We say that (X, Y) -trajectorial uniqueness on \mathbb{E}_0 holds for the equation $\hat{F}_t(X, Y, Z) = 0$ if equalities $\hat{F}_t(\mathbf{x}, \mathbf{y}, \mathbf{z}) = 0$ and $\hat{F}_t(\mathbf{x}, \mathbf{y}', \mathbf{z}') = 0$ for $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \in \mathbb{E}_0$ and $(\mathbf{x}, \mathbf{y}', \mathbf{z}') \in \mathbb{E}_0$ imply that $\mathbf{y} = \mathbf{y}'$. In the next lemma, given \mathbf{x} , we let $G(\mathbf{x})$ denote the set of \mathbf{y} such that $\hat{F}_t(\mathbf{x}, \mathbf{y}, \mathbf{z}) = 0$, $t \in \mathbb{R}_+$, for some \mathbf{z} , where $(\mathbf{x}, \mathbf{y}, \mathbf{z}) \in \mathbb{E}_0$. Note that $G(\mathbf{x}) = \emptyset$ if \mathbf{x} does not belong to the projection of \mathbb{E}_0 on \mathbb{E}^X .

Lemma 3.1. *If (X, Y) -trajectorial uniqueness on \mathbb{E}_0 holds for the equation $\hat{F}_t(X, Y, Z) = 0$, then $\mathbf{\Pi}(Y = \mathbf{y}) = \sup_{\mathbf{x} \in G^{-1}(\mathbf{y})} \mathbf{\Pi}(X = \mathbf{x})$. In particular, $\mathbf{\Pi}(Y = \mathbf{y})$ is specified uniquely.*

Proof. Let U be the subset of Υ such that $F_t(X(v), Y(v), Z(v)) = 0$ and $(X(v), Y(v), Z(v)) \in \mathbb{E}_0$ for $v \in U$. By hypotheses, $\mathbf{\Pi}(\Upsilon \setminus U) = 0$, so $\mathbf{\Pi}(X = \mathbf{x}, Y = \mathbf{y}) = \mathbf{\Pi}(\{X = \mathbf{x}, Y = \mathbf{y}\} \cap U)$. If $v \in U$ and $\mathbf{\Pi}(v) > 0$, then the set $G(X(v))$ contains one element. Hence, either $\mathbf{\Pi}(X = \mathbf{x}, Y = \mathbf{y}) = \mathbf{\Pi}(X = \mathbf{x})$ if $\mathbf{y} \in G(\mathbf{x})$ or $\mathbf{\Pi}(X = \mathbf{x}, Y = \mathbf{y}) = 0$ otherwise. Consequently, $\mathbf{\Pi}(Y = \mathbf{y}) = \sup_{\mathbf{x} \in \mathbb{E}^X} \mathbf{\Pi}(X = \mathbf{x}, Y = \mathbf{y}) = \sup_{\mathbf{x} \in G^{-1}(\mathbf{y})} \mathbf{\Pi}(X = \mathbf{x}) = \sup_{\mathbf{x} \in G^{-1}(\mathbf{y})} \mathbf{\Pi}^X(\mathbf{x})$. \square

Thus, the proof of the theorem is completed by establishing trajectorial uniqueness. In what follows, we implement this programme.

Stochastic equations By (2.1)

$$\begin{aligned} \overline{Q}_k^n(t) &= \overline{Q}_k^n(0) + \sum_{m \in C_k} \left[\int_0^t \mathbf{1} \left(\frac{\overline{Q}_k^n(s-)}{w_{km}} < \min_{\substack{l \in S_m: \\ l \neq k}} \frac{\overline{Q}_l^n(s-)}{w_{lm}} \right) d\overline{A}_m^n(s) \right. \\ &\quad \left. + \int_0^t \mathbf{1} \left(\frac{\overline{Q}_k^n(s-)}{w_{km}} = \min_{\substack{l \in S_m: \\ l \neq k}} \frac{\overline{Q}_l^n(s-)}{w_{lm}} \right) \alpha_{km}^n(ns) d\overline{A}_m^n(s) \right] - \int_0^t \mathbf{1}(\overline{Q}_k^n(s-) > 0) d\overline{B}_k^n(s). \quad (3.1) \end{aligned}$$

We introduce the following latent variables

$$\overline{D}_k^n(t) = \int_0^t \mathbf{1}(\overline{Q}_k^n(s-) > 0) d\overline{B}_k^n(s), \quad (3.2)$$

$$\begin{aligned} \overline{E}_{km}^n(t) &= \int_0^t \mathbf{1}\left(\frac{\overline{Q}_k^n(s-)}{w_{km}} < \min_{\substack{l \in S_m: \\ l \neq k}} \frac{\overline{Q}_l^n(s-)}{w_{lm}}\right) d\overline{A}_m^n(s) \\ &+ \int_0^t \mathbf{1}\left(\frac{\overline{Q}_k^n(s-)}{w_{km}} = \min_{\substack{l \in S_m: \\ l \neq k}} \frac{\overline{Q}_l^n(s-)}{w_{lm}}\right) \alpha_{km}^n(ns) d\overline{A}_m^n(s). \end{aligned} \quad (3.3)$$

They enable us to replace (3.4a) with a system of “continuous” equations. By (3.1), (3.2), and (3.3) for $t \in \mathbb{R}_+$ and $k = 1, 2, \dots, K$

$$\overline{Q}_k^n(t) = \overline{Q}_k^n(0) + \sum_{m \in C_k} \overline{E}_k^n(t) - \overline{D}_k^n(t), \quad (3.4a)$$

$$\int_0^t \overline{Q}_k^n(s-) d\overline{D}_k^n(s) = \int_0^t \overline{Q}_k^n(s-) d\overline{B}_k^n(s), \quad (3.4b)$$

$$\int_0^t \left(\frac{\overline{Q}_k^n(s-)}{w_{km}} - \min_{l \in S_m} \frac{\overline{Q}_l^n(s-)}{w_{lm}} \right) d\overline{E}_{km}^n(s) = 0, \quad (3.4c)$$

$$\overline{A}_m^n(t) = \sum_{k \in S_m} \overline{E}_{km}^n(t). \quad (3.4d)$$

Let $\overline{D}_k^n = (\overline{D}_k^n(t), t \in \mathbb{R}_+)$ and $\overline{E}_{km}^n = (\overline{E}_{km}^n(t), t \in \mathbb{R}_+)$, where $k = 1, 2, \dots, K$ and $m = 1, 2, \dots, M$, $\overline{D}^n = (\overline{D}_k^n, k = 1, 2, \dots, K)$ and $\overline{E}^n = (\overline{E}_{km}^n, k = 1, 2, \dots, K, m = 1, 2, \dots, M)$. The latent variables Z^n are defined by $Z^n = (\overline{E}^n, \overline{D}^n)$ and assume values in $\mathbb{E}^Z = \mathbb{D}(\mathbb{R}_+, \mathbb{R}^{K \times M}) \times \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$. Below, generic elements of spaces \mathbb{E}^X , \mathbb{E}^Y , and \mathbb{E}^Z are denoted as follows: for space \mathbb{E}^X , as $\mathbf{x} = (q_0, \mathbf{a}, \mathbf{b})$, where $q_0 = (q_{0,1}, \dots, q_{0,K}) \in \mathbb{R}^K$, $\mathbf{a} = (\mathbf{a}_m, m = 1, 2, \dots, M) \in \mathbb{D}(\mathbb{R}_+, \mathbb{R}^M)$, and $\mathbf{b} = (\mathbf{b}_k, k = 1, 2, \dots, K) \in \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$, for space \mathbb{E}^Y , as $\mathbf{y} = \mathbf{q}$, where $\mathbf{q} = (\mathbf{q}_k, k = 1, 2, \dots, K) \in \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$, and for space \mathbb{E}^Z , as $\mathbf{z} = (\mathbf{e}, \mathbf{d})$, where $\mathbf{e} = (\mathbf{e}_k, k = 1, 2, \dots, K) \in \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$ and $\mathbf{d} = (\mathbf{d}_k, k = 1, 2, \dots, K) \in \mathbb{D}(\mathbb{R}_+, \mathbb{R}^K)$.

Exponential tightness As pointed out above, the hypotheses of the theorem imply that the X^n LD converge in distribution at rate n to an idempotent variable $X = (Q_0, A, B)$ with idempotent distribution \mathbf{I}^X . Since A and B have continuous paths \mathbf{I}^X -a.e., the sequence (A^n, B^n) is \mathbb{C} -exponentially tight. Let us show that the sequence (A^n, B^n, Y^n, Z^n) is \mathbb{C} -exponentially tight. We repeat the argument of the proof of Lemma 4.1 in Puhalskii [27]. By (3.1), (3.2), and (3.3) the increments of the $\overline{Q}_k^n(t)$, $\overline{E}_{km}^n(t)$ and $\overline{D}_k^n(t)$ are majorised as follows by the increments of $\overline{A}_m^n(t)$ and $\overline{B}_k^n(t)$: for $s < t$

$$|\overline{Q}_k^n(t) - \overline{Q}_k^n(s)| \leq \sum_{m \in C_k} |\overline{A}_m^n(t) - \overline{A}_m^n(s)| + |\overline{B}_k^n(t) - \overline{B}_k^n(s)|, \quad (3.5)$$

$$|\overline{D}_k^n(t) - \overline{D}_k^n(s)| \leq |\overline{B}_k^n(t) - \overline{B}_k^n(s)|, \quad (3.5)$$

$$|\overline{E}_{km}^n(t) - \overline{E}_{km}^n(s)| \leq |\overline{A}_m^n(t) - \overline{A}_m^n(s)|. \quad (3.6)$$

Since the sequences $\{\bar{A}_m^n, n \in \mathbb{N}\}$ and $\{\bar{B}_k^n, n \in \mathbb{N}\}$ are \mathbb{C} -exponentially tight and the sequence $\{\bar{Q}^n(0), n \in \mathbb{N}\}$ is exponentially tight, the above inequalities imply by (A.1) that the sequences $\{\bar{Q}_k^n, n \in \mathbb{N}\}$, $\{\bar{E}_{km}^n, n \in \mathbb{N}\}$ and $\{\bar{D}_k^n, n \in \mathbb{N}\}$ are \mathbb{C} -exponentially tight. Therefore, the sequences Y^n and Z^n are \mathbb{C} -exponentially tight, so the sequence $\{(A^n, B^n, Y^n, Z^n), n \in \mathbb{N}\}$ is \mathbb{C} -exponentially tight. As a consequence, the sequence $\{(X^n, Y^n, Z^n), n \in \mathbb{N}\}$ is exponentially tight.

Large deviation limit Let the (X^n, Y^n, Z^n) LD converge along a subsequence to an idempotent process (X, Y, Z) defined on an idempotent probability space $(\Upsilon, \mathbf{\Pi})$, where $X = (Q_0, A, B)$, $Y = Q$, and $Z = (E, D)$. Our ultimate goal is to show that $\mathbf{\Pi}^Q(\mathbf{q}) = \mathbf{\Pi}(Q = \mathbf{q})$ is specified uniquely. Note that $\mathbf{\Pi}$ -a.e. $A(0) = 0 \in \mathbb{R}^M$, $B(0) = D(0) = 0 \in \mathbb{R}^K$, and $E(0) = 0 \in \mathbb{R}^{K \times M}$. Let us show that $\mathbf{\Pi}$ -a.e. the idempotent processes D_k are absolutely continuous, nonnegative and nondecreasing. Since $\mathbf{P}(\bar{D}_k^n(t) - \bar{D}_k^n(s) < 0) = 0$ for $t \geq s$ and $\liminf_{n \rightarrow \infty} \mathbf{P}(\bar{D}_k^n(t) - \bar{D}_k^n(s) < 0)^{1/n} \geq \mathbf{\Pi}(D_k(t) - D_k(s) < 0)$, we obtain that $\mathbf{\Pi}(D_k \text{ is not nondecreasing}) = \sup_{s < t} \mathbf{\Pi}(D_k(t) - D_k(s) < 0) = 0$, so D_k is nondecreasing $\mathbf{\Pi}$ -a.e. It is therefore nonnegative. Next, by (3.5) $\mathbf{P}(|\bar{D}_k^n(t) - \bar{D}_k^n(s)| > |\bar{B}_k^n(t) - \bar{B}_k^n(s)|) = 0$, so similarly to the above $\mathbf{\Pi}(|D_k(t) - D_k(s)| > |B_k(t) - B_k(s)|) = 0$. Since the function B_k is absolutely continuous $\mathbf{\Pi}$ -a.e., it follows that so is D_k and $\dot{D}_k(t) \leq \dot{B}_k(t)$. A similar argument applied to the E_{km} and using (3.6) shows that the E_{km} are $\mathbf{\Pi}$ -a.e. absolutely continuous and nondecreasing.

Letting $n \rightarrow \infty$ in (3.4a), (3.4b), (3.4c) and (3.4d) we obtain by the continuous mapping principle that $\mathbf{\Pi}$ -a.e.

$$\begin{aligned} Q_k(t) &= Q_{0,k} + \sum_{m \in C_k} E_{km}(t) - D_k(t), \\ \int_0^t Q_k(s) dD_k(s) &= \int_0^t Q_k(s) dB_k(s), \\ \int_0^t \left(\frac{Q_k(s)}{w_{km}} - \min_{l \in S_m} \frac{Q_l(s)}{w_{lm}} \right) dE_{km}(s) &= 0, \\ A_m(t) &= \sum_{k \in S_m} E_{km}(t). \end{aligned}$$

Trajectorial uniqueness We take as \mathbb{E}_0 the subset of $\mathbb{R}^K \times \mathbb{D}(\mathbb{R}_+, \mathbb{R}^{3K+M+KM+1})$ of elements $(q_0, \mathbf{q}, \mathbf{a}, \mathbf{b}, \mathbf{e}, \mathbf{d})$ such that $\mathbf{q}(0) = q_0$, the functions $\mathbf{q}, \mathbf{a}, \mathbf{b}, \mathbf{d}$ and \mathbf{e} are componentwise nonnegative and absolutely continuous, the functions $\mathbf{a}, \mathbf{b}, \mathbf{d}$ and \mathbf{e} are nondecreasing and start at 0, and the following relations hold for $k = 1, 2, \dots, K$, $m = 1, 2, \dots, M$, $t \in \mathbb{R}_+$

$$\dot{\mathbf{q}}_k(t) = \sum_{m \in C_k} \dot{\mathbf{e}}_{km}(t) - \dot{\mathbf{d}}_k(t), \quad \mathbf{q}_k(0) = q_{0,k}, \quad (3.8a)$$

$$\mathbf{q}_k(t)(\dot{\mathbf{d}}_k(t) - \dot{\mathbf{b}}_k(t)) = 0, \quad (3.8b)$$

$$\left(\frac{\mathbf{q}_k(t)}{w_{km}} - \min_{l \in S_m} \frac{\mathbf{q}_l(t)}{w_{lm}} \right) \dot{\mathbf{e}}_{km}(t) = 0, \quad (3.8c)$$

$$\dot{\mathbf{a}}_m(t) = \sum_{k \in S_m} \dot{\mathbf{e}}_{km}(t), \quad \dot{\mathbf{d}}_k(t) \leq \dot{\mathbf{b}}_k(t). \quad (3.8d)$$

We have proved that $\mathbf{\Pi}((Q_0, A, B, Q, E, D) \notin \mathbb{E}_0) = 0$. Hence, $\mathbf{\Pi}(Q = \mathbf{q}) = 0$ if \mathbf{q} either is not absolutely continuous, or componentwise nonnegative, or it is not obtained as a solution of the

latter system of equations. We next prove that this system of equations uniquely specifies \mathbf{q} given q_0 , \mathbf{a} and \mathbf{b} .

Let \mathbf{q}' be another solution, i.e.,

$$\begin{aligned}\dot{\mathbf{q}}'_k(t) &= \sum_{m \in C_k} \dot{\mathbf{e}}'_{km}(t) - \dot{\mathbf{d}}'_k(t), \quad \mathbf{q}'_k(0) = q_{0,k}, \\ \mathbf{q}'_k(t)(\dot{\mathbf{d}}'_k(t) - \dot{\mathbf{b}}_k(t)) &= 0, \\ \left(\frac{\mathbf{q}'_k(t)}{w_{km}} - \min_{l \in S_m} \frac{\mathbf{q}'_l(t)}{w_{lm}}\right) \dot{\mathbf{e}}'_{km}(t) &= 0, \\ \dot{\mathbf{a}}_m(t) &= \sum_{k \in S_m} \dot{\mathbf{e}}'_{km}(t), \quad \dot{\mathbf{d}}'_k(t) \leq \dot{\mathbf{b}}_k(t).\end{aligned}$$

It suffices to prove that

$$\frac{d}{dt} \sum_{k=1}^K |\mathbf{q}_k(t) - \mathbf{q}'_k(t)| \leq 0. \quad (3.9)$$

We have

$$\begin{aligned}\frac{d}{dt} \sum_{k=1}^K |\mathbf{q}_k(t) - \mathbf{q}'_k(t)| &= \sum_{k=1}^K \left((\dot{\mathbf{q}}_k(t) - \dot{\mathbf{q}}'_k(t)) \mathbf{1}(\mathbf{q}_k(t) > \mathbf{q}'_k(t)) \right. \\ &\quad \left. + (\dot{\mathbf{q}}'_k(t) - \dot{\mathbf{q}}_k(t)) \mathbf{1}(\mathbf{q}_k(t) < \mathbf{q}'_k(t)) \right). \quad (3.10)\end{aligned}$$

By (3.8a)

$$\dot{\mathbf{q}}_k(t) - \dot{\mathbf{q}}'_k(t) = \sum_{m \in C_k} (\dot{\mathbf{e}}_{km}(t) - \dot{\mathbf{e}}'_{km}(t)) + (\dot{\mathbf{d}}'_k(t) - \dot{\mathbf{d}}_k(t)). \quad (3.11)$$

If $\mathbf{q}_k(t) > \mathbf{q}'_k(t)$, then $\mathbf{q}_k(t) > 0$, so by (3.8b) $\dot{\mathbf{d}}_k(t) = \dot{\mathbf{b}}_k(t)$, which implies by the inequality $\dot{\mathbf{d}}'_k(t) \leq \dot{\mathbf{b}}_k(t)$ and (3.11) that $\dot{\mathbf{q}}_k(t) - \dot{\mathbf{q}}'_k(t) \leq \sum_{m \in C_k} (\dot{\mathbf{e}}_{km}(t) - \dot{\mathbf{e}}'_{km}(t))$ a.e on the set $\{\mathbf{q}_k(t) > \mathbf{q}'_k(t)\}$. Similarly, $\dot{\mathbf{q}}'_k(t) - \dot{\mathbf{q}}_k(t) \leq \sum_{m \in C_k} (\dot{\mathbf{e}}'_{km}(t) - \dot{\mathbf{e}}_{km}(t))$ a.e on the set $\{\mathbf{q}'_k(t) > \mathbf{q}_k(t)\}$. Hence, by (3.10)

$$\begin{aligned}\frac{d}{dt} \sum_{k=1}^K |\mathbf{q}_k(t) - \mathbf{q}'_k(t)| &\leq \sum_{m=1}^M \sum_{k \in S_m} \left((\dot{\mathbf{e}}_{km}(t) - \dot{\mathbf{e}}'_{km}(t)) \mathbf{1}(\mathbf{q}_k(t) > \mathbf{q}'_k(t)) \right. \\ &\quad \left. + (\dot{\mathbf{e}}'_{km}(t) - \dot{\mathbf{e}}_{km}(t)) \mathbf{1}(\mathbf{q}_k(t) < \mathbf{q}'_k(t)) \right).\end{aligned}$$

We prove that on the right-hand side each sum over S_m is nonpositive.

If $\min_{l \in S_m} \mathbf{q}_l(t)/w_{lm} = \min_{l \in S_m} \mathbf{q}'_l(t)/w_{lm}$, then on the set where $\mathbf{q}_k(t) > \mathbf{q}'_k(t)$ we have that $\mathbf{q}_k(t)/w_{km} > \mathbf{q}'_k(t)/w_{km} \geq \min_{l \in S_m} \mathbf{q}_l(t)/w_{lm}$, so by (3.8c) $\dot{\mathbf{e}}_{km}(t) = 0$. Similarly, $\dot{\mathbf{e}}'_{km}(t) = 0$ a.e. on the set where $\mathbf{q}'_k(t) > \mathbf{q}_k(t)$. The required property follows.

If $\min_{l \in S_m} \mathbf{q}_l(t)/w_{lm} < \min_{l \in S_m} \mathbf{q}'_l(t)/w_{lm}$, then, analogously to the preceding argument, $\dot{\mathbf{e}}_{km}(t) = 0$ a.e. on the set where $\mathbf{q}_k(t) \geq \mathbf{q}'_k(t)$, and the sum in question is not greater than $\sum_{k \in S_m} (\dot{\mathbf{e}}'_{km}(t) - \dot{\mathbf{e}}_{km}(t)) = \dot{\mathbf{a}}_m(t) - \dot{\mathbf{a}}_m(t) = 0$. Similarly, if $\min_{l \in S_m} \mathbf{q}_l(t)/w_{lm} > \min_{l \in S_m} \mathbf{q}'_l(t)/w_{lm}$, then $\dot{\mathbf{e}}'_{km}(t) = 0$ on the set where $\mathbf{q}_k(t) \leq \mathbf{q}'_k(t)$ so that this sum is not greater than $\sum_{k \in S_m} (\dot{\mathbf{e}}_{km}(t) - \dot{\mathbf{e}}'_{km}(t)) = 0$. Inequality (3.9) has been proved.

Evaluating the limit idempotent distribution By Lemma 3.1

$$\mathbf{\Pi}(Q = \mathbf{q}) = \sup_{(q_0, \mathbf{a}, \mathbf{b}) \in \Delta(\mathbf{q})} \mathbf{\Pi}(Q_0 = q_0, A = \mathbf{a}, B = \mathbf{b}), \quad (3.12)$$

where $\Delta(\mathbf{q})$ is the set of $(q_0, \mathbf{a}, \mathbf{b})$ such that (3.8a), (3.8b), (3.8c) and (3.8d) are satisfied.

We conclude the proof by evaluating the right-hand side of (3.12). Let $\mathbf{I}(\mathbf{q}) = -\log \mathbf{\Pi}(Q = \mathbf{q})$. Then

$$\mathbf{I}(\mathbf{q}) = \inf_{(q_0, \mathbf{a}, \mathbf{b}) \in \Delta(\mathbf{q})} (\mathbf{I}^{Q_0}(q_0) + \mathbf{I}^{A, B}(\mathbf{a}, \mathbf{b})) = \mathbf{I}^{Q_0}(\mathbf{q}(0)) + \inf_{(\mathbf{a}, \mathbf{b}) \in \tilde{\Delta}(\mathbf{q})} \int_0^\infty \psi(\dot{\mathbf{a}}(t), \dot{\mathbf{b}}(t)) dt, \quad (3.13)$$

where $\tilde{\Delta}(\mathbf{q})$ is the set of (\mathbf{a}, \mathbf{b}) such that (3.8a), (3.8b), (3.8c) and (3.8d) are satisfied with $\mathbf{q}(0)$ as q_0 . The definition of $L(x, y)$ implies that $\psi(\dot{\mathbf{a}}(t), \dot{\mathbf{b}}(t)) \geq L(\mathbf{q}(t), \dot{\mathbf{q}}(t))$ a.e. for $(\mathbf{a}, \mathbf{b}) \in \tilde{\Delta}(\mathbf{q})$. Hence, $\mathbf{I}(\mathbf{q}) \geq \mathbf{I}^Q(\mathbf{q})$.

In order to prove the reverse inequality, we assume that $\mathbf{I}^Q(\mathbf{q}) < \infty$, so $L(\mathbf{q}(t), \dot{\mathbf{q}}(t)) < \infty$ a.e. Since the infimum in the definition (2.3) of $L(x, y)$ is attained if finite, the sets $\Gamma(t), t \in \mathbb{R}_+$, consisting of $(a, b) \in N(\mathbf{q}(t), \dot{\mathbf{q}}(t))$ such that $\psi(a, b) = L(\mathbf{q}(t), \dot{\mathbf{q}}(t))$ are nonempty for almost all t . Also, the graph $\{(t, (a, b)) \in \mathbb{R}_+ \times \mathbb{R}^{M+K} : (a, b) \in \Gamma(t)\}$ is measurable with respect to the product of the Lebesgue σ -algebra on \mathbb{R}_+ and Borel σ -algebra on \mathbb{R}^{M+K} . Therefore, by a measurable selection theorem (see, e.g., Clark [8]) there exist Lebesgue measurable functions $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{b}}$ such that $(\tilde{\mathbf{a}}(t), \tilde{\mathbf{b}}(t)) \in \Gamma(t)$ a.e. Letting $\mathbf{a}(t) = \int_0^t \tilde{\mathbf{a}}(s) ds$ and $\mathbf{b}(t) = \int_0^t \tilde{\mathbf{b}}(s) ds$, we obtain that $(\mathbf{a}, \mathbf{b}) \in \tilde{\Delta}(\mathbf{q})$ and $\int_0^\infty \psi(\dot{\mathbf{a}}(t), \dot{\mathbf{b}}(t)) dt = \int_0^\infty L(\mathbf{q}(t), \dot{\mathbf{q}}(t)) dt$. Thus the infimum on the right of (3.13) is equal to $\int_0^\infty L(\mathbf{q}(t), \dot{\mathbf{q}}(t)) dt$, hence, $\mathbf{I}(\mathbf{q}) = \mathbf{I}^Q(\mathbf{q})$ and $\mathbf{\Pi}^Q(\mathbf{q}) = \exp(-\mathbf{I}^Q(\mathbf{q}))$.

Acknowledgements The first author is thankful to Yuri Suhov for drawing his attention to the join-the-shortest-queue setting.

A Review of idempotent probability and large deviation convergence

Let Υ be a set. A function $\mathbf{\Pi}$ from the power set of Υ to $[0, 1]$ is called an idempotent probability if $\mathbf{\Pi}(\Gamma) = \sup_{v \in \Gamma} \mathbf{\Pi}(\{v\})$, $\Gamma \subset \Upsilon$ and $\mathbf{\Pi}(\Upsilon) = 1$. The pair $(\Upsilon, \mathbf{\Pi})$ is called an idempotent probability space. For economy of notation, we denote $\mathbf{\Pi}(v) = \mathbf{\Pi}(\{v\})$. A property $\mathcal{P}(v)$, $v \in \Upsilon$, pertaining to the elements of Υ is said to hold $\mathbf{\Pi}$ -a.e. if $\mathbf{\Pi}(\mathcal{P}(v))$ does not hold $= 0$. A function f from a set Υ equipped with idempotent probability $\mathbf{\Pi}$ to a set Υ' is called an idempotent variable. The idempotent distribution of an idempotent variable f is defined as the set function $\mathbf{\Pi} \circ f^{-1}(\Gamma) = \mathbf{\Pi}(f \in \Gamma)$, $\Gamma \subset \Upsilon'$. If f is the canonical idempotent variable that is defined by $f(v) = v$, then it has $\mathbf{\Pi}$ as the idempotent distribution. Υ' -valued idempotent variables f and f' are said to be independent if $\mathbf{\Pi}(f = v', f' = v'') = \mathbf{\Pi}(f = v')\mathbf{\Pi}(f' = v'')$ for all $v', v'' \in \Upsilon'$. Independence of finite collections of idempotent variables is defined similarly. A collection $(X_t, t \in \mathbb{R}_+)$ of \mathbb{R}^ℓ -valued idempotent variables on Υ , where $\ell \in \mathbb{N}$, is called an idempotent process. The functions $(X_t(v), t \in \mathbb{R}_+)$ for various $v \in \Upsilon$ are called trajectories (or paths) of X . Idempotent processes are said to be independent if they are independent as idempotent variables with values in the associated function space.

If Υ is, in addition, a metric space and the sets $\{v \in \Upsilon : \mathbf{\Pi}(v) \geq \alpha\}$ are compact for all $\alpha \in (0, 1]$, then $\mathbf{\Pi}$ is called a deviability. Obviously, $\mathbf{\Pi}$ is a deviability if and only if $\mathbf{I}(v) = -\log \mathbf{\Pi}(\{v\})$ is an

action functional. If f is a continuous mapping from Υ to another metric space Υ' , then $\mathbf{\Pi} \circ f^{-1}$ is a deviability on Υ' . As a matter of fact, for the latter property to hold, one can only require that f be continuous on the sets $\{v \in \Upsilon : \mathbf{\Pi}(v) \geq \alpha\}$ for $\alpha \in (0, 1]$. In general, f is said to be a Luzin idempotent variable if $\mathbf{\Pi} \circ f^{-1}$ is a deviability on Υ' .

Let $\{\mathbf{P}_n, n \in \mathbb{N}\}$ be a sequence of probability measures on a metric space \mathbb{E} endowed with Borel σ -algebra and let $\mathbf{\Pi}$ be a deviability on \mathbb{E} . Let $m_n \rightarrow \infty$ as $n \rightarrow \infty$. The sequence $\{\mathbf{P}_n, n \in \mathbb{N}\}$ is said to large deviation converge (LD converge) at rate m_n to $\mathbf{\Pi}$ as $n \rightarrow \infty$ if $\lim_{n \rightarrow \infty} \left(\int_{\mathbb{E}} f(z)^{m_n} \mathbf{P}_n(dz) \right)^{1/m_n} = \sup_{z \in \mathbb{E}} f(z)\mathbf{\Pi}(z)$ for every bounded continuous \mathbb{R}_+ -valued function f on \mathbb{E} . This definition is equivalent to requiring that the inequalities $\limsup_{n \rightarrow \infty} \mathbf{P}_n(F)^{1/m_n} \leq \mathbf{\Pi}(F)$ and $\liminf_{n \rightarrow \infty} \mathbf{P}_n(G)^{1/m_n} \geq \mathbf{\Pi}(G)$ hold for all closed sets F and all open sets G , respectively. Therefore, the sequence $\{\mathbf{P}_n, n \in \mathbb{N}\}$ LD converges at rate m_n to $\mathbf{\Pi}$ if and only if it obeys the LDP for scale m_n with action functional $\mathbf{I}(z) = -\log \mathbf{\Pi}(z)$. We favour the term ‘‘LD convergence’’ over ‘‘the LDP’’ as being more natural for our approach. The deviability $\mathbf{\Pi}$ is said to be an LD limit point of the \mathbf{P}_n for rate n if each subsequence $\{\mathbf{P}_{n_k}, k \in \mathbb{N}\}$ of $\{\mathbf{P}_n, n \in \mathbb{N}\}$ contains a further subsequence $\{\mathbf{P}_{n_{k_l}}, l \in \mathbb{N}\}$ that LD converges to $\mathbf{\Pi}$ at rate n_{k_l} as $l \rightarrow \infty$. The sequence $\{\mathbf{P}_n, n \in \mathbb{N}\}$ is said to be exponentially tight on order m_n if for arbitrary $\epsilon > 0$ there exists a compact subset E of \mathbb{E} such that $\limsup_{n \rightarrow \infty} \mathbf{P}_n(\mathbb{E} \setminus E)^{1/m_n} < \epsilon$. An exponentially tight sequence possesses LD limit points. Thus, one can prove LD convergence of the \mathbf{P}_n by proving that exponential tightness holds and that there is a unique LD limit point. We will say that a sequence of random variables with values in a metric space is exponentially tight if so is the sequence of their laws.

LD convergence of probability measures can be also expressed as LD convergence in distribution of the associated random variables to idempotent variables. In the setting of stochastic processes, this point of view enables one to consider the LD limit as a dynamical system rather than as ‘‘a mass function’’ on the space of trajectories. We say that a sequence $\{X_n, n \in \mathbb{N}\}$ of random variables defined on probability spaces $(\Omega_n, \mathcal{F}_n, \mathbf{P}_n)$ and assuming values in \mathbb{E} LD converges in distribution at rate m_n as $n \rightarrow \infty$ to a Luzin idempotent variable X defined on an idempotent probability space $(\Upsilon, \mathbf{\Pi})$ and assuming values in \mathbb{E} if the sequence of the probability laws of the X_n LD converges to the idempotent distribution of X at rate m_n . Conversely, LD convergence of a sequence $\{\mathbf{P}_n, n \in \mathbb{N}\}$ of probability measures on \mathbb{E} to a deviability $\mathbf{\Pi}$ on \mathbb{E} can be expressed as LD convergence in distribution if one considers the canonical setting. The continuous mapping principle, known as the contraction principle for the LDP, states that if the X_n LD converge in distribution to X and f is a continuous function from \mathbb{E} to another metric space, then the $f(X_n)$ LD converge in distribution to $f(X)$. We will use the extension, also referred to as the continuous mapping principle, where the function f is allowed to be a measurable function that is continuous only a.e. with respect to the idempotent distribution of X . For a detailed discussion, see Garcia [16]. The definition of a limit point for LD convergence in distribution is similar to that for LD convergence of probability measures.

If the X_n are stochastic processes with trajectories in a Skorohod space $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^\ell)$, then the sequence X_n is said to be \mathbb{C} -exponentially tight on order m_n if it is exponentially tight on order m_n and each LD limit point $\mathbf{\Pi}$ of the distributions of the X_n is an idempotent distribution of a continuous-path idempotent process in the sense that $\mathbf{\Pi}(\mathbb{D}(\mathbb{R}_+, \mathbb{R}^\ell) \setminus \mathbb{C}(\mathbb{R}_+, \mathbb{R}^\ell)) = 0$. The sequence

X_n is \mathbb{C} -exponentially tight on order m_n if and only if

$$\begin{aligned} & \lim_{L \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbf{P}(|X_n(0)| > L)^{1/m_n} = 0, \\ & \lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbf{P} \left(\sup_{\substack{s, t \in [0, T]: \\ |s-t| \leq \delta}} |X_n(t) - X_n(s)| > \epsilon \right)^{1/m_n} = 0, \quad T \in \mathbb{R}_+, \epsilon > 0. \end{aligned} \quad (\text{A.1})$$

If the sequence X_n is \mathbb{C} -exponentially tight, then a limit point of the X_n for LD convergence in distribution may be considered as an idempotent process with trajectories in $\mathbb{C}(\mathbb{R}_+, \mathbb{R}^\ell)$. Conversely, if the X_n LD converge in distribution to a continuous-path idempotent process, then the sequence X_n is \mathbb{C} -exponentially tight.

We also note that, as a consequence of the continuous mapping principle, if a sequence of stochastic processes (X_n, Y_n) assuming values in \mathbb{R}^{ℓ_1} and \mathbb{R}^{ℓ_2} , respectively, LD converges in distribution in $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^{\ell_1}) \times \mathbb{D}(\mathbb{R}_+, \mathbb{R}^{\ell_2})$ to a continuous-path idempotent process, then the LD convergence also holds in $\mathbb{D}(\mathbb{R}_+, \mathbb{R}^{\ell_1} \times \mathbb{R}^{\ell_2})$.

References

- [1] M. Alanyali and B. Hajek. On large deviations in load sharing networks. *Ann. Appl. Prob.*, 8:67–97, 1998.
- [2] M. Alanyali and B. Hajek. On large deviations of Markov processes with discontinuous statistics. *Ann. Appl. Prob.*, 8:45–66, 1998.
- [3] R. Atar and P. Dupuis. Large deviations and queueing networks: methods for rate function identification. *Stochastic Process. Appl.*, 84(2):255–296, 1999.
- [4] V. M. Blinovskii and R. L. Dobrushin. Process level large deviations for a class of piecewise homogeneous random walks. In *The Dynkin Festschrift: Markov Processes and their Applications*, pages 1–59. Birkhäuser, Boston, 1994.
- [5] A. A. Borovkov and A. A. Mogul'skiĭ. Large deviations for Markov chains in the positive quadrant. *Uspekhi Mat. Nauk*, 56(5(341)):3–116, 2001.
- [6] M. Boué, P. Dupuis, and R. S. Ellis. Large deviations for small noise diffusions with discontinuous statistics. *Probab. Theory Related Fields*, 116(1):125–149, 2000.
- [7] H. Chen and A. Mandelbaum. Discrete flow networks: bottleneck analysis and fluid approximations. *Math. Oper. Res.*, 16(2):408–446, 1991.
- [8] F.H. Clarke. *Optimization and Nonsmooth Analysis*. Wiley, 1983.
- [9] P. Dupuis and R. Ellis. Large deviations for Markov processes with discontinuous statistics. II. *Probab. Theory Related Fields*, 91:153–194, 1992.
- [10] P. Dupuis and R. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley, 1997.
- [11] P. Dupuis and R. S. Ellis. The large deviation principle for a general class of queueing systems, I. *Trans. Amer. Math. Soc.*, 347:2689–2751, 1995.

- [12] P. Dupuis, R. S. Ellis, and A. Weiss. Large deviations for Markov processes with discontinuous statistics. I. *Ann. Prob.*, 19:1280–1297, 1991.
- [13] S. N. Ethier and T. G. Kurtz. *Markov Processes. Characterization and Convergence*. Wiley, 1986.
- [14] P. J. Fleming and B. Simon. Heavy traffic approximations for a system of infinite servers with load balancing. *Prob. Engrg. Inform. Sci.*, 13:251–273, 1999.
- [15] R. D. Foley and D. R. McDonald. Join the shortest queue: stability and exact asymptotics. *Ann. Appl. Prob.*, 11(3):569–607, 2001.
- [16] J. Garcia. An extension of the contraction principle. *J. Theoret. Probab.*, 17(2):403–434, 2004.
- [17] I. Ignatiouk-Robert. Large deviations of Jackson networks. *Ann. Appl. Prob.*, 10(3):962–1001, 2000.
- [18] I. Ignatiouk-Robert. Large deviations for processes with discontinuous statistics. *Ann. Probab.*, 33(4):1479–1508, 2005.
- [19] J. Jacod and A.N. Shiryaev. *Limit Theorems for Stochastic Processes*. Springer, 1987.
- [20] A. P. Korostelëv and S. L. Leonov. An action functional for a diffusion process with discontinuous drift. *Theory Probab. Appl.*, 37(3):543–550, 1992. (in Russian: *Teor. Veroyatn. i Primen.*, 1992, v. 37, no. 3, pp. 570-576).
- [21] A. P. Korostelëv and S. L. Leonov. Action functional for diffusions in discontinuous media. *Probab. Theory Related Fields*, 94(3):317–333, 1993.
- [22] R. Sh. Liptser and A. N. Shiryaev. *Theory of Martingales*. Kluwer, 1989.
- [23] K. Majewski. Large deviation bounds for single class queueing networks and their calculation. *Queueing Syst.*, 48(1-2):103–134, 2004.
- [24] A. Puhalskii. On functional principle of large deviations. In V.Sazonov and T.Shervashidze, editors, *New Trends in Probability and Statistics*, volume 1, pages 198–218. VSP/Moks’las, 1991.
- [25] A. Puhalskii. The method of stochastic exponentials for large deviations. *Stochastic Process. Appl.*, 54(1):45–70, 1994.
- [26] A. Puhalskii. *Large Deviations and Idempotent Probability*. Chapman & Hall/CRC, 2001.
- [27] A.A. Puhalskii. The action functional for the Jackson network. (submitted), 2005.
- [28] K. Ramanan and P. Dupuis. Large deviation properties of data streams that share a buffer. *Ann. Appl. Probab.*, 8(4):1070–1129, 1998.
- [29] A. Ridder and A. Shwartz. Large deviations without principle: Join the shortest queue. In *Mathematical Methods in Operations Research*. 2005.
- [30] A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis*. Chapman & Hall, 1995.
- [31] S. R. E. Turner. Large deviations for Join the Shorter Queue. In *Analysis of communication networks: call centres, traffic and performance (Toronto, ON, 1998)*, volume 28 of *Fields Inst. Commun.*, pages 95–108. Amer. Math. Soc., Providence, RI, 2000.